Web Results **1 - 10** of about **36,000** for **modifying database storage through logs in write ahead logging sy**

Scholarly articles for **modifying database storage through logs in write ahead logging system**

Using write protected data structures to improve ... - by Sullivan - 13 citations
A transactional memory service in an extensible ... - by Saito - 4 citations
Lightweight transactions on networks of workstations - by Papathanasiou - 4 citations

Contents
Setting Up Your Backup **System** and **Storage** ... Configuring the **Database** Server to Read **Ahead** ... **Modifying Database Logging** Mode with ON-Monitor (UNIX) ...
publib.boulder.ibm.com/infocenter/ ids9help/topic/com.ibm.admin.doc/admin02.htm - 109k - Cached - Similar pages

Lightweight Transactions on Networks of Workstations
Some examples of **systems** that use the **Write-Ahead Logging** Protocols are RVM ... can be found in the remote undo **logs** are copied back to the remote **database**, ...
www.ics.forth.gr/carv/r-d-activities/ networkMem/icdcs98/online.html - 52k - Cached - Similar pages

Citations: Reimplementing the Cedar File **System** Using **Logging** and ...
Write ahead logging systems [4, 6, 12, 32] accumulate small updates in a **log** and ... Improving **Storage System** Availability with D-GRAID - Muthian Sivathanu ...
citeseer.ist.psu.edu/context/8260/0 - 37k - Cached - Similar pages

[PDF] Microsoft PowerPoint - ch17recovery
File Format: PDF/Adobe Acrobat - View as HTML
Modifying the **database** without ensuring that the transaction will ... stable **storage**. (This rule is called the **write-ahead logging** or WAL ...
www.cs.brown.edu/courses/ cs127/lectures/L20.recovery.pdf - Similar pages

Database Management **Systems**
These **systems** persistently **store** the sorts of object and pointer ... PostgreSQL implemented **write-ahead-logging** (which is what he is talking about) with ...
philip.greenspun.com/panda/**databases**-choosing - 78k - Cached - Similar pages

UC Berkeley Operating **System** Prelims - Spring 1999
One key idea in implementing ACID transactions is to use **write-ahead logging**. The **log** is maintained on stable **storage**, and all modifications are first ...
guir.berkeley.edu/projects/ osprelims/summaries/transactions.html - 25k - Cached - Similar pages

[PS] A Transactional Memory Service in an Extensible Operating **System** ...
File Format: Adobe PostScript - View as Text
Efficient buffering must comply with the **write ahead logging** (WAL) rule, ... **System** Call Function **storage** = trans open(path) Opens the **database** file path ...
www.cs.washington.edu/homes/bershad/Papers/trans.ps - Similar pages

PostgreSQL

%Create dump tool for **write-ahead logs** for use in determining transaction ... Currently non-global **system** tables must be in the default **database** tablespace. ...
www.postgresql.org/docs/faqs.TODO.html - 54k - Sep 5, 2005 - Cached - Similar pages

[DOC] Web **Store** IT Whitepaper
File Format: Microsoft Word 2000 - View as HTML
At the **database** level, the Web **Storage System** supports transaction **logging**. It uses **write-ahead** transaction **logs** to ensure data integrity **through** redundancy ...
www.somorita.com/exchange2000/webstore_white_paper.doc - Sep 5, 2005 - Similar pages

[PDF] Failure-Atomic File Access in an Interposed Network **Storage System**
File Format: PDF/Adobe Acrobat - View as HTML
proach is quite similar to **write-ahead logging** that might. be taken on a journaling central ... **system** protocol, without **modifying** the client or server. ...
www.cs.duke.edu/ari/publications/slice-recovery.pdf - Similar pages

Goooooooooogle ▶
Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

modifying database storage through Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

Google

Web Images Groups News Froogle Local more »

write ahead logging and warm cache Search Advanced Search Preferences

The "AND" operator is unnecessary -- we include all search terms by default. [details]

Web Results **1 - 10** of about **49,600** for **write ahead logging and warm cache**. (**0.36** seconds)

[PDF] Integrating reliable memory in databases
File Format: PDF/Adobe Acrobat - View as HTML
straints imposed by **write-ahead logging** [Gray78]. To order ... it allows the database to start with a **warm cache** in the event ...
www.eecs.umich.edu/CoVirt/ papers/databaseReliableMemJournal.pdf - Similar pages

[PDF] Integrating Reliable Memory in Databases
File Format: PDF/Adobe Acrobat - View as HTML
warm cache) [Sullivan93, Elhardt84, Bhide93]. Storing the **log** and/or the buffer **cache** in reliable ... straints imposed by **write-ahead logging** [Gray78]. To ...
www.eecs.umich.edu/CoVirt/ papers/databaseReliableMem.pdf - Similar pages
[More results from www.eecs.umich.edu]

[PS] TeXPS: dvi->PostScript Driver dvitps, Version 2.19 of June 23 ...
File Format: Adobe PostScript - View as Text
The Episode **warm cache** read rate is a bit slower than the corresponding JFS rate, ... Locking and Partial Rollbacks using **Write-ahead Logging**. ...
tedanderson.home.mindspring.com/episode.ps.gz - Similar pages

PostgreSQL
Auto-fill the free space map by scanning the buffer **cache** or by checking pages ... **Write-Ahead Log**. Eliminate need to **write** full pages to WAL before page ...
www.postgresql.org/docs/faqs.TODO.html - 54k - Sep 5, 2005 - Cached - Similar pages

[PS] A Transactional Memory Service in an Extensible Operating System ...
File Format: Adobe PostScript - View as Text
Efficient buffering must comply with the **write ahead logging** (WAL) rule, ... Buffer **cache** was **warm**. 0 5 10 15 20 25 30. T1 T2a T2c. Traversal type ...
www.cs.washington.edu/homes/bershad/Papers/trans.ps - Similar pages

[PS] A Comparative Study of **Log**-Only and In-Place UpdateBased Temporal ...
File Format: Adobe PostScript - View as Text
Updating can be done in-place, with **write-ahead logging**. ... be in the **cache** on the disk, and can be retrieved without an additional disk seek operation. ...
www.idi.ntnu.no/grupper/db/ research/tech_papers/CIKM2000/CIKM2000.ps - Similar pages

[PS] Advanced Systems Topics Part I of III Steven Hand Lent Term 2005 6 ...
File Format: Adobe PostScript - View as Text
use **write-ahead log** - recall transactions from Part IB CSAA. Advanced System Topics -- Database Storage ... Overall: approx local FS (if **warm cache**. . .) ...
www.cl.cam.ac.uk/Teaching/current/AdvSysTop/ast-smh.ps - Similar pages

[PDF] Advanced Systems Topics Part I of III
File Format: PDF/Adobe Acrobat - View as HTML
NTFS directly supports [infinite] **write-ahead log**. – Some extensibility research in this area (eg ... Overall: approx local FS (if **warm cache**. . .) ...
www.cl.cam.ac.uk/Teaching/current/AdvSysTop/ast-smh.pdf - Similar pages

Integrating Coherency and Recoverability in Distributed Systems ...
At each commit point, after the **log** records have been **written** to the storage ...
page copy (cold **cache**) | 171.9 | 43 | | page copy (**warm cache**) | 57.8 | 135 ...
www.usenix.org/publications/ library/proceedings/osdi/full_papers/feeley.a - Similar pages

[PS] Integrating Coherency and Recoverability in Distributed Systems ...
File Format: Adobe PostScript
Log-based coherency is an extension to **write-ahead** redo **logging**, ...
(_sec/page) (MBytes/s) page copy (cold **cache**) 171.9 43 page copy (**warm cache**) 57.8 135 ...
www.usenix.org/publications/library/ proceedings/osdi/full_papers/feeley.ps - Similar pages

Goooooooooogle ▶
Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

write ahead logging and warm cache Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

Google Web Images Groups News Froogle Local more »
write ahead logging Search Advanced Search Preferences

Web Results **1 - 10** of about **650,000** for **write ahead logging**. (**0.32** seconds)

Write-Ahead Logging (WAL)
Write-Ahead Logging (WAL) is a standard approach to transaction **logging**.
Its detailed description may be found in most (if not all) books about ...
www.physiol.ox.ac.uk/Computing/ Online_Documentation/postgresql/wal.html - 6k - Cached - Similar pages

Write-Ahead Logging (WAL) - SQL Database Reference Material ...
SQL.org aims to be both a portal to SQL resources on the internet, and a source of original SQL-related content.
www.sql.org/sql-database/postgresql/manual/wal.html - 14k - Cached - Similar pages

Derby **Write Ahead Log** Format
This document describes the storage format of Derby **Write Ahead Log**. ... A good description of **Write Ahead Logging**, and how a **log** is typically implemented, ...
db.apache.org/derby/papers/logformats.html - 31k - Sep 4, 2005 - Cached - Similar pages

Write-Ahead Logging (WAL)
Write-Ahead Logging (WAL). Table of Contents; 12.1. General Description ...
Write Ahead Logging (WAL) is a standard approach to transaction **logging**. ...
jamesthornton.com/postgres/7.3/postgres/wal.html - 15k - Cached - Similar pages

Write-Ahead Logging (WAL)
Write-Ahead Logging (WAL). Table of Contents; 11.1. General Description ...
Write Ahead Logging (WAL) is a standard approach to transaction **logging**. ...
jamesthornton.com/postgres/7.2/postgres/wal.html - 15k - Cached - Similar pages
[More results from jamesthornton.com]

Write-Ahead Logging (WAL)
Write Ahead Logging (WAL) is a standard approach to transaction **logging**.
Its detailed description may be found in most books about transaction processing. ...
www.redhat.com/docs/manuals/database/ RHDB-7.1.3-Manual/admin_user/wal.html - 8k - Cached - Similar pages

Write ahead logging -- Facts, Info, and Encyclopedia article
Write ahead logging. [Categories: Data management] In (The branch of engineering science that studies (with the aid of computers) computable processes and ...
www.absoluteastronomy.com/ encyclopedia/w/wr/**write_ahead_logging**.htm - 15k - Cached - Similar pages

Write Ahead Log - What does **Write Ahead Log** stand for? Definitions ...
What does **Write Ahead Log** stand for? Definition of **Write Ahead Log** in the list of acronyms provided by the Free Online Dictionary and Thesaurus.
acronyms.thefreedictionary.com/Write+Ahead+Log - 20k - Cached - Similar pages

[PDF] ARIES/NT: A Recovery Method Based on **Write-Ahead Logging** for ...
File Format: PDF/Adobe Acrobat - View as HTML
write-ahead logging (WAL) is the. method of choice in most commercial systems because ... Using **Write-Ahead. Logging**, IBM. Re-. search Report RJ6846, ...
www.vldb.org/conf/1989/P337.PDF - Similar pages

[PDF] Course Project – Step Two
File Format: PDF/Adobe Acrobat - View as HTML
actions, **write-ahead logging**, and forward and restart processing. ... **log** (arWAL).
The **write-ahead log** is a file and should reuse your file management code ...
hssl.cs.jhu.edu/~randal/416/proj2.pdf - Similar pages

Goooooooooogle ▶
Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

write ahead logging Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

Google Web Images Groups News Froogle Local more »
disaster recovery with write ahead logging sys Search Advanced Search Preferences

Web Results **1 - 10** of about **205,000** for **disaster recovery with write ahead logging system**. (**0.24** seconds)

Scholarly articles for **disaster recovery with write ahead logging system**

Algorithms for the Management of Remote Backup Data ... - by Mohan - 29 citations
An efficient scheme for providing high availability - by Bhide - 14 citations
Coordinator log transaction execution protocol - by Stamos - 26 citations

ICDE 1993: 511-518
... Considerations in Replicated Database **Systems** for **Disaster** Protection. ... A Concurrency Control and **Recovery** Method Using **Write-Ahead Logging** for ...
www.informatik.uni-trier.de/ ~ley/db/conf/icde/MohanTO93.html - 19k - Sep 4, 2005 - Cached - Similar pages

ICDE 1991: 718-727
... A Concurrency Control and **Recovery** Method Using **Write-Ahead Logging** for Linear ... and Data Integrity in Transaction **Systems** Using **Write-Ahead Logging**. ...
www.informatik.uni-trier.de/ ~ley/db/conf/icde/Mohan91.html - 25k - Sep 4, 2005 - Cached - Similar pages
[More results from www.informatik.uni-trier.de]

[PDF] Microsoft PowerPoint - 5A-CR
File Format: PDF/Adobe Acrobat - View as HTML
Use the **written** information to undo. 4. **Write**. **Write--Ahead Logging**. **Ahead Logging** ... **Disaster Recovery**. **Disaster Recovery**. **System** pairs for availability ...
www.cc.gatech.edu/classes/ AY2005/cs4803enc_fall/papers/5A-CR.pdf - Similar pages

Network Appliance - SnapManager 3.0 for Microsoft Exchange ...
The transaction-**logging** mechanism uses a **write-ahead** algorithm, ensuring that every ... and rapid **recovery** in the event of a **system** failure or **disaster**. ...
www.netapp.com/tech_library/3198.html - 48k - Cached - Similar pages

Network Appliance - Microsoft Exchange 5.5, NetApp Filers and ...
The ESE uses "**write ahead**" transaction **logging** for both the Directory Store and ... to deliver highly available and **disaster recovery** solutions with add-on ...
www.netapp.com/tech_library/3120.html - 51k - Cached - Similar pages

SQL 2000 Backup and Restore. Advice on protecting your database ...
Overall backup strategy; **Write ahead log** files; **Log** files and restore; Backup methods ... 2) Backup your databases or backup the **system** databases ...
www.computerperformance.co.uk/ SQL2000/SQL_backup_restore.htm - 22k - Cached - Similar pages

Exchange 2003 - Circular **Logging** of Storage Groups
Exchange relies on transaction or **write-ahead** logs to store events before they ... The fatal flaw with Circular **Logging** is it restricts **disaster recovery**. ...
www.computerperformance.co.uk/ exchange2003/exchange2003_circular_**logging**.htm - 19k - Cached - Similar pages

[PDF] Repeating History Beyond ARIES Repeating History Beyond ARIES
File Format: PDF/Adobe Acrobat - View as HTML
Using **write-ahead logging**, efficient, fine granule locking. Using **write-ahead logging**, efficient, fine granule ... Remote site back up and **disaster recovery** ...
www.dcs.napier.ac.uk/~vldb99/ IndustrialSpeakerSlides/vldb99_aries_slides.pdf - Similar pages

DiSC - Repeating History Beyond ARIES
... Analysis of **Recovery** in a Database **System** Using a **Write-Ahead Log** Protocol. ... for the Management of Remote Backup Data Bases for **Disaster Recovery**. ...
www.sigmod.org/sigmod/disc/p_repeatinghistorc.htm - 43k - Cached - Similar pages

Disaster Recovery in SharePoint Products and Technologies
Writing a **Disaster Recovery** Plan, **Writing** a **Disaster Recovery** Plan ... If you want to use the SQL **Log** Shipping feature as a **disaster recovery** method for ...
www.microsoft.com/technet/ prodtechnol/sppt/reskit/c2861881x.mspx - 101k - Sep 4, 2005 - Cached - Similar pages

Gooooooooogle ▶
Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

disaster recovery with write ahead lc Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

Google Home - Advertising Programs - Business Solutions - About Google

- Papers on ARIES
- Patents on ARIES

Books which Cover ARIES

1. Weihl, W. Transaction Processing Techniques, Chapter 13 in **Distributed Systems**, S. Mullender (Ed.), 2nd Edition, ISBN 0-201-62427-3, ACM Press, 1993.
2. Kumar, V. (Ed.). **Performance of Concurrency Control Mechanisms in Centralized Database Systems**, ISBN 0-13-065442-6, Prentice Hall, 1995.
3. Ramamritham, K., Chrysanthis, P. **Advances in Concurrency Control and Transaction Processing**, ISBN 0-81-867405-9, IEEE Computer Society Press, 1996.
4. Franklin, M. Concurrency Control and Recovery, In **The Handbook of Computer Science and Engineering**, A. Tucker (Ed.), ISBN 0-849-32909-4, CRC Press, 1997.
5. Stonebraker, M., Hellerstein, J. (Eds.). **Readings in Database Systems**, 3rd Edition, ISBN 1-558-60523-1, Morgan Kaufmann Publishers, 1998. Slides on ARIES.
6. Abdelguerfi, M., Wong, K.-F. **Parallel Database Techniques**, ISBN 0-8186-8398-8, IEEE Computer Science Press, July 1998.
7. Kumar, V., Hsu, M. (Eds.). **Recovery Mechanisms In Database Systems**, ISBN 0-13-614215-X, Prentice Hall, 1998.
8. Hvasshovd, S.-O. **Recovery in Parallel Database Systems**, 2nd Edition, ISBN 3-528-15411-X, Vieweg Verlag, 1999.
9. Ramakrishnan, R., Gehrke, J. **Database Management Systems**, 2nd Edition, ISBN 0-07-232206-3, McGraw-Hill, August 1999. Slides on ARIES.
10. Saake, G., Heuer, A. **Datenbanken: Implementierungstechniken**, ISBN 3-8266-0513-6, MITP-Verlag, May 1999. German slides on ARIES.
11. Elmasri, R., Navathe, S. **Fundamentals of Database Systems**, 3rd Edition, ISBN 0-201-74153-9, Addison Wesley, 2000.
12. Silberschatz, A., Korth, H., Sudarshan, S. **Database System Concepts**, 4th Edition, ISBN 0-07-228363-7, McGraw-Hill, July 2001. Slides on ARIES.
13. Weikum, G., Vossen, G. **Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control and Recovery**, ISBN 1-555860-508-8, Morgan Kaufmann, 2001.
14. Albano, A. **Costruire Sistemi per Basi di Dati**, ISBN 88-7192-106-2, Addison Wesley

Longman, Milano, 2001.

List of Courses or Qualifying Examinations in which ARIES is Included

Universities in the following 20 countries are teaching ARIES: **Australia, Canada, Denmark, England, Finland, France, Germany, Greece, India, Iran, Israel, Italy, Korea, New Zealand, Norway, Singapore, Spain, Sweden, Taiwan, USA.**

1. Baylor University, Waco (Greg Speegle)
2. Ben-Gurion University of the Negev, Israel (Ehud Gudes)
3. Brandeis University (Liuba Shrira)
4. California State University at Chico (Renee Renner)
5. Carnegie-Mellon University (Anastassia Ailamaki)
6. Columbia University (Luis Gravano)
7. Cornell University (Johannes Gehrke, Jayavel Shanmugasundaram), another
8. Duke University (Jeff Chase), another
9. George Mason University
10. Georgia Institute of Technology
11. Harvard University (Margo Seltzer), another (Dan Ellard)
12. Hebrew University of Jerusalem, Israel (Catriel Beeri, Yehoshua Sagiv)
13. Helsinki University of Technology, Finland (Eljas Soisalon-Soininen)
14. H.K.B.K. College of Engineering, India
15. Illinois Institute of Technology (Ratko Orlandic)
16. Indian Institute of Science, Bangalore, India (Jayant Haritsa)
17. Indian Institute of Technology at Bombay, India (S. Sudarshan), another (Krithi Ramamritham, Sunita Sarawagi)
18. Iowa State University
19. James Cook University, Australia (Hossein Ghodosi)
20. Johns Hopkins University, Baltimore (Randal Burns)
21. Karlstad University, Sweden (Niklas Nikitin)
22. Korea Advanced Institute of Science and Technology, Korea (Kyu-Young Whang)
23. Kyungpook National University, Korea (Young-Chul Park)
24. L'Ecole Nationale Supérieure des Télécommunications (ENST), France (Luc Bouganim)
25. Linkopings Universitet, Sweden (Olof Johansson)
26. Massey University, New Zealand (Markus Kirchberg)
27. National Dong Hwa University, Taiwan (Shiow-yang Wu)
28. National Technical University of Athens, Greece (Timos Sellis)
29. National University of Singapore, Singapore (Lee Mong Li)
30. North Dakota State University (Victor Shi)
31. Northeastern University, Boston (Betty Salzberg)
32. Norwegian University of Science and Technology, Trondheim, Norway (Roger Midtstraum), another
33. Oregon Graduate Institute (Lois Delcambre)
34. Osmania University, Hyderabad, India
35. Pennsylvania State University (Thomas Keefe)
36. Pohang University of Science and Technology, Pohang, Korea (In Jun Choi)

37. Princeton University (Andrea LaPaugh)
38. Purdue University (K.C.L. Van Zandt), another (Sunil Prabhakar)
39. Rensselaer Polytechnic Institute (Sibel Adali)
40. San Francisco State University (Marguerite Murphy)
41. Seoul National University, Korea
42. Sharif University of Technology, Iran (Rasool Jalili)
43. Siena College (Scott Vandenberg)
44. Simon Fraser University, Vancouver, Canada (Ke Wang)
45. Stanford University
46. Technion - Israel Institute of Technology, Haifa, Israel (Oded Shmueli)
47. Universidad de Puerto Rico, Puerto Rico (Manuel Rodríguez Martínez)
48. Università di Pisa, Pisa, Italy (Antonio Albano)
49. Università di Roma "La Sapienza", Rome, Italy (Tiziana Catarci)
50. Universitat Autònoma de Barcelona, Barcelona, Spain (Xavier Binefa, Fernando Vilariño)
51. Universität des Saarlandes, Saarbruecken, Germany (Gerhard Weikum)
52. Universität Leipzig, Leipzig, Germany (Erhard Rahm)
53. Universität Mannheim, Mannheim, Germany (Mila Majster-Cederbaum)
54. Universität Trier, Trier, Germany (Michael Ley)
55. University at Buffalo (Jan Chomicki)
56. University of Alberta, Edmonton, Canada (Osmar Zaiane)
57. University of Athens, Athens, Greece (Yannis Ioannidis)
58. University of Birmingham, Birmingham, England (Alan Sexton)
59. University of British Columbia, Vancouver, Canada (Ed Knorr)
60. University of California at Berkeley (Eric Brewer, Joe Hellerstein), Qualifier
61. University of California at Irvine (Sharad Mehrotra), Qualifier
62. University of California at Santa Cruz (Randal Burns)
63. University of Canberra, Canberra, Australia (Ric Jentzsch)
64. University of Canterbury, New Zealand (Tanja Mitrovic, Bruce Blum)
65. University of Central Oklahoma (Lester McCann)
66. University of Copenhagen, Copenhagen, Denmark (Lars G. T. Jørgensen)
67. University of Helsinki, Helsinki, Finland (Seppo Sippu)
68. University of Illinois at Urbana-Champaign (Chad Peiper)
69. University of Iowa (Ramon Lawrence)
70. University of Ioannina, Greece (Evaggelia Pitoura)
71. University of Kentucky (Alexander Dekhtyar), qualifier
72. University of Manitoba, Canada (Sajid Hussain)
73. University of Maryland (Michael Franklin)
74. University of Massachusetts at Lowell (John Sieg)
75. University of Memphis (King-Ip (David) Lin)
76. University of Michigan (Jignesh Patel)
77. University of Minnesota at Duluth (Rich Maclin)
78. University of Minnesota at Minneapolis (Shashi Shekhar)
79. University of Missouri - Columbia (Frederick Springsteel)
80. University of New South Wales, Australia (Mohammad Nabil)
81. University of Passau, Germany (Donald Kossman)
82. University of Patras, Greece
83. University of Pennsylvania (Susan Davidson)
84. University of Skovde, Sweden (Jonas Mellin)
85. University of South Carolina at Columbia (Caroline Eastman)
86. University of Texas at Arlington (Sharma Chakravarthy)
87. University of Texas at Austin (Don Batory)

88. University of Texas at El Paso (Karen Ward)
89. University of Texas at San Antonio (Weining Zhang)
90. University of Vermont, Burlington (Byung Lee)
91. University of Waterloo, Canada (Tamer Ozsu), another
92. University of Western Ontario, Canada (Sylvia Osborn)
93. University of Wisconsin at Madison (David DeWitt)
94. Wellesley College (Scott Anderson)

Systems which have Incorporated ARIES Technology

- DB2 Universal Database for z/OS and OS/390 (previously called DB2/MVS)
- DB2 Universal Database for Linux, Windows, AIX, Solaris, HP-UX, NUMA-Q, OS/2 (originally called OS/2 Extended Edition Database Manager)
- DB2 Server for VM and VSE (previously called SQL/DS)
- Cloudscape
- Lotus Domino (French document), see also [1] [2]
- MQSeries persistent messaging and queuing system (previously called Message Queue Manager)
- Tivoli Storage Manager (TSM, previously called AdStar Distributed Storage Manager (ADSM))
- Encina recoverable file system
- VM Shared File System
- Starburst extensible DBMS
- QuickSilver distributed operating system
- Microsoft SQL Server and NT file system
- Sybase ASE 11.9.2
- Clustra DBMS
- University of Wisconsin's
  - Gamma Database Machine
  - Shore persistent object system
  - EXODUS extensible DBMS
  - Paradise GIS system
  - Minirel DBMS
  - Minibase DBMS
- Cornell University's
  - PREDATOR object-relational DBMS
  - Mars ARIES simulator (see also)
- Seoul National University's
  - Soprano client-server object storage system
  - SNU RDBMS Platform (SRP)
- University of Pisa's Java Relational System (JRS)
- University of Melbourne's Aditi deductive database system (see also)
- Australian National University's Platypus Object Store
- Johns Hopkins University's ARIES Project

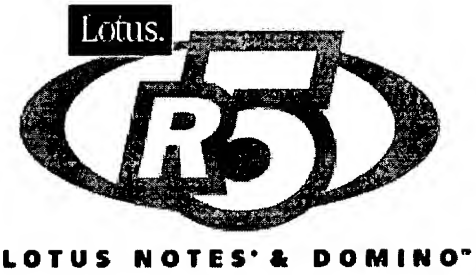

Quotes about ARIES

RECENT ADVANCES IN TRANSACTION MANAGEMENT
Invited Speaker: Prof. Alan Fekete (University of Sydney)

ADC '93, Fourth Australian Database Conference
1-2 February 1993, Brisbane, Queensland, Australia

Summary: For the past decade, transaction management had a reputation as a field where research proj irrelevant to practitioners, since all systems used the same long-understood techniques (namely two-ph write-ahead logging). This has now changed.

In the past few years, there have been several exciting advances in transaction management that seem influence future commercial systems. One is the invention and publication of improved techniques for of transaction management. In particular, an exciting series of papers have come from the ARIES proj Mohan at IBM Almaden Research Laboratory. There are new algorithms which provide concurrency c indices, recovery compatible with fine-grained locking, and concurrency control allowing long-runnin advance has been driven by advanced applications such as distributed programming and collaborative domains the traditional transaction model (with short, sequential, isolated transactions) is inadequate. I transaction models have been proposed, such as sagas and nested transactions. Each new model needs for managing concurrency and failure.

This talk will present some of these exciting new ideas. A unifying theme will be the identification of between different aspects of transaction management.

The speaker: Alan Fekete holds a PhD from Harvard University. In 1987-88 he worked at the Laborat Science at Massachusetts Institute of Technology. Since 1988 he has been at the University of Sydney, teaches the Database Systems course. His research interest is the theory of distributed systems, especia algorithms for transaction management in distributed database management systems. He has been invc project to understand concurrency control for nested transaction systems; this work is presented in the "Atomic Transactions" by Lynch, Merritt, Weihl and Fekete (published by Morgan Kaufmann).

Charles Bash's Review of the Main ARIES Article: Mohan, C., Haderle, D., Lindsay, B., Pir Schwarz, P. *ARIES: A Transaction Recovery Method Supporting Fine-Granularity Locking and Partia Using Write-Ahead Logging*, **ACM Transactions on Database Systems**, Vol. 17, No. 1, March 1992,

New algorithms for database recovery and rollbacks are described. The paper assumes that the databas ahead logging (WAL), but it describes in fine detail how the various activities during the update, rollb phases are to act so as to maximize concurrency and minimize both overhead and time. In their introdu authors also provide an excellent description of the current state of the art of logging, failures, and rec

The paper is broken into 12 sections and has an extensive bibliography (101 citations). The sections ar introduction, goals, an overview of ARIES, a description of the major data structures, a discussion of t are part of normal processing (including transaction failure), a description of restart processing (after s description of the impact of checkpoints during restart, the methods necessary for media recovery, top (independent transactions kicked off by running transactions such as file extension), recovery paradign problems caused by them), properties of other WAL-based methods (including references to several c implementations), and a summary of the attributes of ARIES.

This paper is excellent both for those who wish to know more about restart/recovery methods and for t

to improve them. My only problem reading the paper was the profusion of three-letter acronyms. Thes reduce the length of the paper significantly, and they are defined well on first usage. Due to the length however, the reader sometimes wishes to find that definition again, which may be difficult. A short glc help, and if it included a reference to the defining section, the reader could again find the details easily

I would like to thank the authors for documenting this excellent work, which clearly will improve the s recovery/restart. I recommend it for all who are involved in database management system design or wi the internals better.

Reminiscences on Influential Papers, ACM SIGMOD Record, September 1998

Prof. Betty Salzberg, Northeastern University, salzberg@ccs.neu.edu

[C. Mohan, D. Haderle, B. Lindsay, H. Pirahesh and P. Schwarz, "ARIES: A Transaction Recovery M Fine-Granularity Locking and Partial Rollbacks Using Write-Ahead Logging," ACM Transactions on Systems, 17(1):94-162, March 1992.]

The ARIES paper was important for me because it enabled me to envision the mechanisms of recovery systems clearly. For example, I saw how Log Sequence Numbers (LSNs) are used to enforce Write-Ah (WAL). WAL says that before a page with an update on it made by an uncommitted transaction can be (overwriting the previous version of the page), the pre-image of the updated record must be on disk so But it is always important to understand some mechanism by which a theoretical rule can be enforced. commonly used for WAL is the LSN. The LSN = L on a database page P in the buffer in main memory the log record of the most recent update on P . Log records contain preimages of updated records and l written sequentially in increasing LSN order. If the LSN of the most recent log record written to disk is WAL implies that P cannot yet be written to disk. First, a portion of the log containing the log record v must be written to disk. This is one of many recovery mechanisms I did not know and I think many oth researchers did not know until preprints of the ARIES paper were made available. Reading the ARIES influenced much of my subsequent research. My research on concurrency and recovery for B-link-tree methods (the II-tree and the hB-II tree) for example, uses LSNs to determine whether an index page ha since the last visit. (If it has not been updated, a new search through the tree can be avoided.) Issues of and support for fine-granularity locking, exposed in the ARIES paper, were essential in the II-tree and The concept of page-oriented vs. logical UNDO was explained in the ARIES paper and used in the II-t tree. My research on transactional workflow (DSDT) and on online reorganization uses the method of from log records (from ARIES) to recreate system tables and/or reestablish the state of an ongoing app is almost impossible for me to imagine thinking of a database system without ARIES style recovery.

in Edinburgh on Tuesday 7th September 1999.

Reminiscences on Influential Papers, ACM SIGMOD Record, March 2000

Prof. Pat O'Neil, UMass/Boston, poneil@cs.umb.edu

[C. Mohan. Concurrency Control and Recovery Methods for B+-Tree Indexes: ARIES/KVL and ARIE Performance of Concurrency Control Mechanisms in Centralized Database Systems, Prentice-Hall 19S 065442-6]

For a number of years I was uncomfortable about my understanding of how locking is used by vendors transactional phantoms, especially since I knew that the "Predicate Locking" approach mentioned in m been dropped by System R many years ago (see: Astrahan et al., TODS 1(2), 1976). When the KVL ar approaches used at IBM were first published by Mohan in VLDB-90 and SIGMOD-92, I was anxious thoroughly, but I found it difficult to fully grasp the concepts in a quick reading. It wasn't until a few y began covering the combined KVL and IM paper above in detail, and presenting it to my database inte now believe that the ideas underlying these locking protocols are probably the most subtle in the datab they are not easily grasped, and since all the researchers I know are extremely busy, I think they have r attention than they deserve. I have heard practitioners complain jokingly that Mohan's papers seem des the detail necessary for experienced programmers to perform immediate implementation. There is a ce truth to this, and I for one find it a wonderful thing. I think an excellent database internals text could b simply expanding on the ideas in this paper (latches, locking by hashing, lock durations, logging, B-tre etc.). The text would be particularly valuable in an academic setting because the techniques covered, s were a revelation to me, are ones that are ACTUALLY USED by IBM database programmers. I canno that many university researchers (both faculty and students) could overcome perceived isolation from i realities by studying this seminal work. Even practicing database system programmers, inside and outs have not already spent time on this paper, would be well advised to expand their horizons by reading it

Papers on ARIES

1. Rothermel, K., Mohan, C. *ARIES/NT: A Recovery Method Based on Write-Ahead Logging for Nested Transactions*, **Proc. 15th International Conference on Very Large Data Bases**, Amsterdam, August 1989. A longer version of this paper is available as **IBM Research Report RJ6650**, IBM Almaden Research Center, January 1989. Received the "10 Year Best Impact Paper Award" at VLDB99. Abstract Citations DBLP ACM ResearchIndex

2. Mohan, C., Narang, I., Palmer, J. *A Case Study of Problems in Migrating to Distributed Computing: Page Recovery Using Multiple Logs in the Shared Disks Environment*, **IBM Research Report RJ7343**, IBM Almaden Research Center, March 1990. Abstract Citations ResearchIndex

3. Mohan, C. *Commit_LSN: A Novel and Simple Method for Reducing Locking and Latching in Transaction Processing Systems*, **Proc. 16th International Conference on Very Large Data Bases**, Brisbane, August 1990. Also available as **IBM Research Report RJ7344**, IBM Almaden Research Center, February 1990. A slightly revised version appears in **Performance of Concurrency Control Mechanisms in Centralized Database Systems**, V. Kumar (Ed.),

Prentice Hall, 1995. Abstract Citations DBLP ACM ResearchIndex

4. Mohan, C. *ARIES/KVL: A Key-Value Locking Method for Concurrency Control of Multiaction Transactions Operating on B-Tree Indexes*, **Proc. 16th International Conference on Very Large Data Bases**, Brisbane, August 1990, pp392-405. A different version of this paper is available as IBM Research Report RJ7008, IBM Almaden Research Center, September 1989. Abstract Citations DBLP ACM ResearchIndex

5. Mohan, C., Pirahesh, H. *ARIES-RRH: Restricted Repeating of History in the ARIES Transaction Recovery Method*, **Proc. 7th International Conference on Data Engineering**, Kobe, April 1991. Also available as **IBM Research Report RJ7342**, IBM Almaden Research Center, July 1990. Abstract Citations DBLP ResearchIndex

6. Mohan, C., Narang, I. *Recovery and Coherency-Control Protocols for Fast Intersystem Page Transfer and Fine-Granularity Locking in a Shared Disks Transaction Environment*, **Proc. 17th International Conference on Very Large Data Bases**, Barcelona, September 1991. A longer version of this paper is available as **IBM Research Report RJ8017**, IBM Almaden Research Center, March 1991. Abstract Citations DBLP ResearchIndex

7. Mohan, C., Narang, I., Silen, S. *Solutions to Hot Spot Problems in a Shared Disks Transaction Environment*, **Proc. 4th International Workshop on High Performance Transaction Systems**, Asilomar, September 1991. Also available as **IBM Research Report RJ8281**, IBM Almaden Research Center, August 1991. Abstract Citations ResearchIndex

8. Mohan, C., Haderle, D., Lindsay, B., Pirahesh, H., Schwarz, P. *ARIES: A Transaction Recovery Method Supporting Fine-Granularity Locking and Partial Rollbacks Using Write-Ahead Logging*, **ACM Transactions on Database Systems**, Vol. 17, No. 1, March 1992, pp94-162. Reprinted in **Readings in Database Systems**, 3rd Edition, M. Stonebraker, J. Hellerstein (Eds.), Morgan Kaufmann Publishers, 1998. Reprinted in **Recovery Mechanisms In Database Systems**, V. Kumar, M. Hsu (Eds.), Prentice Hall, 1998. Also available as **IBM Research Report RJ6649**, IBM Almaden Research Center, January 1989; Revised November 1990. Abstract Citations DBLP ACM ResearchIndex

9. Mohan, C., Narang, I. *Efficient Locking and Caching of Data in the Multisystem Shared Disks Transaction Environment*, **Proc. 3rd International Conference on Extending Database Technology**, Vienna, March 1992. Also available as **IBM Research Report RJ8301**, IBM Almaden Research Center, August 1991. Abstract Citations DBLP ResearchIndex

10. Mohan, C., Levine, F. *ARIES/IM: An Efficient and High Concurrency Index Management Method Using Write-Ahead Logging*, **Proc. ACM SIGMOD International Conference on Management of Data**, San Diego, June 1992. A longer version of this paper is available as **IBM Research Report RJ6846**, IBM Almaden Research Center, August 1989. Abstract Citations DBLP ACM ResearchIndex

11. Mohan, C., Narang, I. *Algorithms for Creating Indexes for Very Large Tables Without Quiescing Updates*, **Proc. ACM SIGMOD International Conference on Management of Data**, San Diego, June 1992. A longer version of this paper is available as **IBM Research Report RJ8016**, IBM Almaden Research Center, March 1991. Abstract Citations DBLP ACM ResearchIndex

12. Mohan, C., Pirahesh, H., Lorie, R. *Efficient and Flexible Methods for Transient Versioning of*

Records to Avoid Locking by Read-Only Transactions, **Proc. ACM SIGMOD International Conference on Management of Data**, San Diego, June 1992. Also available as **IBM Research Report RJ8683**, IBM Almaden Research Center, March 1992. Abstract Citations DBLP ACM ResearchIndex

13. Mohan, C., Narang, I. *Data Base Recovery in Shared Disks and Client-Server Architectures*, **Proc. 12th International Conference on Distributed Computing Systems**, Yokohama, June 1992. Also available as **IBM Research Report RJ8685**, IBM Almaden Research Center, March 1992. Abstract Citations ResearchIndex

14. Mohan, C., Treiber, K., Obermarck, R. *Algorithms for the Management of Remote Backup Data Bases for Disaster Recovery*, **Proc. 9th International Conference on Data Engineering**, Vienna, April 1993. A longer version of this paper is available as **IBM Research Report RJ7885**, IBM Almaden Research Center, December 1990; Revised June 1991. Abstract Citations DBLP ResearchIndex

15. Mohan, C. *ARIES/LHS: A Concurrency Control and Recovery Method Using Write-Ahead Logging for Linear Hashing with Separators*, **Proc. 9th International Conference on Data Engineering**, Vienna, April 1993. A longer version of this paper is available as **IBM Research Report RJ8682**, IBM Almaden Research Center, March 1992. Abstract Citations DBLP ResearchIndex

16. Mohan, C., Narang, I. *An Efficient and Flexible Method for Archiving a Data Base*, **Proc. ACM SIGMOD International Conference on Management of Data**, Washington, D.C., May 1993. A *corrected* version of this paper is available as **IBM Research Report RJ9733**, IBM Almaden Research Center, March 1993. DBLP ACM ResearchIndex

17. Mohan, C. *A Cost-Effective Method for Providing Improved Data Availability* During *DBMS Restart Recovery After a Failure*, **Proc. 19th International Conference on Very Large Data Bases**, Dublin, August 1993. Also available as **IBM Research Report RJ8114**, IBM Almaden Research Center, May 1991. DBLP ResearchIndex

18. Mohan, C., Haderle, D. *Algorithms for Flexible Space Management in Transaction Systems Supporting Fine-Granularity Locking*, **Proc. 4th International Conference on Extending Database Technology**, Cambridge, March 1994. A longer version of this paper is available as **IBM Research Report RJ9732**, IBM Almaden Research Center, March 1994. DBLP ResearchIndex

19. Mohan, C., Dievendorff, R. *Recent Work on Distributed Commit Protocols, and Recoverable Messaging and Queuing*, **Data Engineering**, Vol. 17, No. 1, March 1994. DBLP ResearchIndex

20. Mohan, C., Narang, I. *ARIES/CSA: A Method for Database Recovery in Client-Server Architectures*, **Proc. ACM SIGMOD International Conference on Management of Data**, Minneapolis, May 1994. Also available as **IBM Research Report RJ9742**, IBM Almaden Research Center, March 1994. DBLP ACM ResearchIndex

21. Mohan, C., Narang, I. *Locking and Latching Techniques for Transaction Processing Systems Supporting the Shared Disks Architecture*, **Research Report**, IBM Almaden Research Center, October 1994.

22. Rane, S., Seshadri, S., Mohan, C. *Concurrency Control and Recovery Algorithms for hcC-trees*, **IBM Research Report**, IBM Almaden Research Center, February 1995.

23. Mohan, C. *Disk Read-Write Optimizations and Data Integrity in Transaction Systems Using Write-Ahead Logging*, **Proc. 11th International Conference on Data Engineering**, Taipei, March 1995. Also available as **IBM Research Report RJ9741**, IBM Almaden Research Center, March 1994.

24. Mohan, C. *Concurrency Control and Recovery Methods for B^+-Tree Indexes: ARIES/KVL and ARIES/IM*, **Performance of Concurrency Control Mechanisms in Centralized Database Systems**, V. Kumar (Ed.), Prentice Hall, 1995. Also available as **IBM Research Report RJ9715**, IBM Almaden Research Center, March 1994.

25. Choy, D., Mohan, C. *Locking Protocols for Two-Tier Indexing of Partitioned Data*, **Proc. International Workshop on Advanced Transaction Models and Architectures**, Goa, August-September 1996.

26. Kornacker, M., Mohan, C., Hellerstein, J. *Concurrency and Recovery in Generalized Search Trees*, **Proc. ACM SIGMOD International Conference on Management of Data**, Tucson, May 1997.

27. Josten, J., Mohan, C., Narang, I., Teng, J. *DB2's Use of the Coupling Facility for Data Sharing*, **IBM Systems Journal**, Vol. 36, No. 2, 1997.

28. Mohan, C. *Repeating History Beyond ARIES*, **Invited paper** for receiving 10 Year Best Impact Paper Award, **Proc. 25th International Conference on Very Large Data Bases**, Edinburgh, September 1999. Slides of keynote speech. Abstract

29. Mohan, C., Barber, R., Watts, S., Somani, A., Zaharioudakis, M. *Evolution of Groupware for Business Applications: A Database Perspective on Lotus Domino/Notes*, **Proc. 26th International Conference on Very Large Databases**, Cairo, September 2000. Abstract DBLP

30. Narang, I., Mohan, C., Brannon, K., Subramanian, M. *Coordinated Backup and Recovery between Database Management Systems and File Systems*, Submitted for Publication, IBM Almaden Research Center, October 2001. Abstract

31. Mohan, C. *An Efficient Method for Performing Record Deletions and Updates Using Index Scans*, **Proc. 28th International Conference on Very Large Databases**, Hong Kong, August 2002. Abstract

---

Patents on ARIES

1. Levine, F., Mohan, C. *Method for Concurrent Record Access, Insertion, Deletion and Alteration Using an Index Tree*, **United States Patent 4,914,569**, IBM, April 1990. **Taiwan Patent NI-34575**, February 1990. **Canada Patent 1,285,072**, June 1991. **Korea Patent 0,052,225**, June 1992. **Republic of China Patent 0,027,768**, July 1994. **France Patent 0,314,292**, April 1996. **Germany Patent 3,855,213,208**, April 1996.

This method (*ARIES/IM*) has been implemented in DB2 Common Server. Some of the ideas have also been implemented in SQL/DS and the VM Shared File System. With enhancements, ARIES/IM has been implemented in DB2/MVS V4.

2. Levine, F., Mohan, C. *Method and Apparatus for Concurrent Modification of an Index Tree in a Transaction Processing System Utilizing Selective Indication of Structural Modification Operations*, **United States Patent 5,123,104**, IBM, June 1992. **Sri Lanka Patent 0,010,014**, November 1989. **Taiwan Patent NI-40987**, December 1990. **Republic of China Patent 0,022,452**, February 1993. **Philippines Patent 0,027,313**, May 1993. **Korea Patent 0,063,350**, July 1993. **Thailand Patent 0,003,973**, September 1994. **European (France, Germany, Italy, Netherlands, Spain, Sweden, Switzerland) Patent 0,336,035**, November 1995. **HO Patent 712/1996**, April 1996. **Japan Patent 2,505,040**, April 1996.

   This method (*ARIES/IM*) has been implemented in DB2 Common Server. With enhancements, ARIES/IM has been implemented in DB2/MVS V4.

3. Mohan, C., Obermarck, R., Treiber, K. *Concurrently Applying Redo Records to Backup Database in a Log Sequence Using Single Queue Server per Queue at a Time*, **United States Patent 5,170,480**, IBM, December 1992. **Japan Patent 1,868,704**, September 1994.

   This method has been implemented as part of the Remote Site Recovery (RSR) feature of IMS/ESA V5.

4. Mohan, C. *Transaction Processing System and Method With Reduced Locking*, **United States Patent 5,247,672**, IBM, September 1993. **Japan Patent 1,938,731**, June 1995. **European Patent 0,442,715**, December 1997.

   This method (*Commit_LSN*) has been implemented in DB2/MVS V3 and V4 (Commit_LSN feature is called *Lock Avoidance*), and MQSeries/MVS (Message Queue Manager/ESA).

5. Mohan, C., Narang, I. *Non-Blocking Serialization for Caching Data in a Shared Cache*, **United States Patent 5,276,835**, IBM, January 1994. **Japan Patent 2,059,253**, June 1996.

   This method is part of the S/390 Parallel Sysplex Coupling Facility. It is exploited by DB2/MVS V4.

6. Mohan, C., Narang, I., Teng, J. *Method for Managing Database Recovery from Failure of a Shared Store in a System Including a Plurality of Transaction-Based Systems of the Write-Ahead Logging Type*, **United States Patent 5,280,611**, IBM, January 1994. **European Patent 0,541,381**, July 1997.

   This method is part of the S/390 Parallel Sysplex Coupling Facility and DB2/MVS V4.

7. Lorie, R., Mohan, C., Pirahesh, H. *Multiple Version Database Concurrency Control System*, **United States Patent 5,280,612**, IBM, January 1994.

8. Mohan, C., Narang, I. *Non-Blocking Serialization for Removing Data from a Shared Cache*, **United States Patent 5,287,473**, IBM, February 1994. **Japan Patent 2,505,939**, April 1996.

   This method is part of the S/390 Parallel Sysplex Coupling Facility. It is exploited by DB2/MVS

V4.

9. Haderle, D., Lindsay, B., Mohan, C., Pirahesh, H., Schwarz, P. *Method for Managing Subpage Concurrency Control and Partial Transaction Rollback in a Transaction-Oriented System of the Write-Ahead Logging Type*, **United Kingdom Patent 0,295,424**, IBM, April 1994. **France Patent 0,295,424**, April 1994. **Germany Patent 3,889,254,508**, April 1994.

   This method (*ARIES*) has been implemented in DB2/MVS, DB2 Common Server, Workstation Data Save Facility/VM (WDSF/VM), ADSTAR Distributed Storage Manager (ADSM), Message Queue Manager/ESA (MQSeries), Starburst extensible DBMS, QuickSilver distributed operating system, Transarc's Encina product suite, and University of Wisconsin's Gamma and EXODUS DBMSs, and SHORE persistent object system. It has also been implemented in Microsoft's SQL Server and NT File System.

10. Mohan, C., Narang, I. *Fast Intersystem Page Transfer in a Data Sharing Environment with Record Locking*, **United States Patent 5,327,556**, IBM, July 1994.

    Some of the algorithms in this patent form the basis of the methods which support the shared disks architecture in DB2/MVS V4.

11. Mohan, C. *Method for Providing Data Availability in a Transaction-Oriented System During Restart After a Failure*, **United States Patent 5,333,303**, IBM, July 1994.

12. Josten, J., Masatani, T., Mohan, C., Narang, I., Teng, J. *Efficient Data Base Access Using a Shared Electronic Store in a Multi-System Environment with Shared Disks*, **United States Patent 5,408,653**, IBM, April 1995.

    This method has been implemented in DB2/MVS V4.

13. Mohan, C. *A Method and Means for Detecting Partial Page Writes and Avoiding Initializing New Pages on DASD in a Transaction Management System Environment*, **United States Patent 5,418,940**, IBM, May 1995.

    This method has been implemented in the ADSTAR Distributed Storage Manager (ADSM) and DB2 Common Server.

14. Bhide, A., Copeland, G., Goyal, A., Hsiao, H.-I, Jhingran, A., Mohan, C. *Asynchronous Replica Management in Shared Nothing Architectures*, **United States Patent 5,440,727**, IBM, August 1995.

15. Dievendorff, R., Mohan, C. *System and Method for Storing Persistent and Non-Persistent Queued Data and for Recovering the Persistent Data Responsive to a System Restart*, **United States Patent 5,452,430**, IBM, September 1995. **UK Patent 0,623, 877**, January 1999.

    This method has been implemented in Message Queue Manager/ESA (MQSeries).

16. Mohan, C., Narang, I., Teng, J. *Partial Page Write Detection for a Shared Cache Using a Bit Pattern Written at the Beginning and End of Each Page*, **United States Patent 5,455,942**, IBM, October 1995.

    This method is part of the S/390 Parallel Sysplex Coupling Facility. It is exploited by DB2/MVS

V4.

17. Haderle, D., Mohan, C. *Method for Managing Logging and Locking of Page Free Space Information in a Transaction Processing System*, **United States Patent 5,455,944**, IBM, October 1995.

    This method has been implemented in DB2/MVS V3.

18. Mohan, C., Narang, I. *Method and Means for Archiving Modifiable Pages in a Log Based Transaction Management System*, **United States Patent 5,455,946**, IBM, October 1995. **Japan Patent 2,505,112**, April 1996.

    This method has been implemented in the ADSTAR Distributed Storage Manager (ADSM) and in DB2/MVS V4.

19. Dievendorff, R., Mohan, C. *Fault-Tolerant Transaction-Oriented Data Processing*, **United States Patent 5,465,328**, IBM, November 1995.

    This method has been implemented in Message Queue Manager/ESA (MQSeries).

20. Elko, D., Frey, J., Mohan, C., Narang, I., Nick, J., Strickland, J., Swanson, M. *Multiple Processor System having Software for Selecting Shared Cache Entries of an Associated Castout Class for Transfer to an DASD with One I/O Operation*, **United States Patent 5,493,668**, IBM, February 1996. **Japan Patent 2,765,672**, April 1998.

    This method is part of the S/390 Parallel Sysplex Coupling Facility.

21. Elko, D., Frey, J., Isenberg, J., Mohan, C., Narang, I., Nick, J., Strickland, J., Swanson, M. *Sysplex Shared Data Coherency Method*, **United States Patent 5,537,574**, IBM, July 1996. **Japan Patent 2,837,785**, October 1998.

    This method is part of the S/390 Parallel Sysplex Coupling Facility.

22. Choy, D., Mohan, C. *Multi-Tiered Indexing Method for Partitioned Data*, **United States Patent 5,551,027**, IBM, August 1996.

23. Mohan, C., Narang, I. *Method for Non-Hierarchical Lock Management in a Multi-System Shared Data Environment*, **United States Patent 5,551,046**, IBM, August 1996.

    This method has been implemented in DB2/MVS V4.

24. Mohan, C. *System and Method for Performing Record Deletions Using Index Scans*, **United States Patent 6,009,425**, IBM, December 1999.

    This method has been implemented in DB2 UDB for Unix, Windows and OS/2.

25. Barber, R., Herbert, D., Mohan, C., Somani, A., Watts, S., Zaharioudakis, M. *Data Recovery in a Transactional Database Using Write-Ahead Logging and File Caching*, **United States Patent 6,173,292**, IBM, January 2001.

    This method has been implemented in Lotus Domino/Notes R5.

Last updated on 5 August 2004. **C. Mohan, mohan@almaden.ibm.com**